

PATENT APPLICATION TRANSMITTAL LETTER

(Large Entity)

Docket No.

85773-161

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Transmitted herewith for filing under 35 U.S.C. 111 and 37 C.F.R. 1.53 is the patent application of:

Stephen Douglas PETERS et al.

TO: METHOD AND APPARATUS PROVIDING HYPOTHESIS DRIVEN SPEECH MODELLING FOR USE IN

SPEECH RECOGNITION

Enclosed are:

- Certificate of Mailing with Express Mail Mailing Label No.
 8 sheets of drawings.
 A certified copy of a application.
 Declaration Signed. Unsigned.
 Power of Attorney
 Information Disclosure Statement
 Preliminary Amendment
 Other: Assignment (X) Signed

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	36	- 20 =	16	x \$18.00	\$288.00
Indep. Claims	8	- 3 =	5	x \$78.00	\$390.00
Multiple Dependent Claims (check if applicable)	<input type="checkbox"/>				\$0.00
				BASIC FEE	\$760.00
				TOTAL FILING FEE	\$1,438.00

- A check in the amount of \$1,478.00 to cover the filing fee is enclosed.
 The Commissioner is hereby authorized to charge and credit Deposit Account No. 19-2550 as described below. A duplicate copy of this sheet is enclosed.
- Charge the amount of _____ as filing fee.
 Credit any overpayment.
 Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
 Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Dated: December 20, 1999



Signature

Stephan P. Georgiev, Reg. 37,563
 SMART & BIGGAR
 1000 de la Gauchetiere Street West
 Suite 3400
 Montreal, Quebec
 H3B 4W5
 CANADA

*includes \$40.00 assignment registration fee.

cc:

**Title: Method and Apparatus Providing Hypothesis
Driven Speech Modelling for use in Speech
Recognition**

5 **Field of the Invention**

This invention relates to the field of speech recognition and more particularly to a method and apparatus for providing model adaptation on a basis of a spoken utterance. This invention is applicable to speech recognition systems of the types used in directory assistance systems, voice dictation systems and other voice activated services.

15 **Background of the Invention**

In a typical speech recognition application, the user inputs into an input device such as a microphone or telephone set. If valid speech is detected, the speech recognition layer is invoked in an attempt to recognize the unknown utterance. In a commonly used approach, on a first pass search, a fast match algorithm is used to select the top N orthography groups from a speech recognition dictionary. In a second pass the individual orthographies from the selected groups are re-scored using computations on more precise speech models. The top orthography in each of the top two groups is then processed by a rejection algorithm that evaluates if they are sufficiently distinctive from one another so the top choice candidate can be considered to be a valid recognition.

Speech recognition systems can be assigned to two distinct categories, namely speaker-specific (or speaker dependent) and speaker independent. These categories differ 5 primarily in the manner these systems are trained and used.

Training of a speech recognition system establishes a reference memory and speech models to which are assigned speech labels. For speaker-independent systems, training is 10 performed by collecting samples from a large pool of users. For a given speech-recognition task, a speaker-specific (SS) system generally performs better than a speaker-independent (SI) system. Typically, for a speaker-independent system, a single speech model is used for all speakers while in a speaker-specific system, each user is assigned a respective 15 speech model set. Speaker-specific systems are trained by collecting speech samples from the end user. For example, a voice dictation system where a user speaks and the device translates his words into text will most likely be trained 20 by the end user (speaker-specific) since this training fashion can achieve a higher recognition performance. In the event that someone else than the original user wants to use the same device, that device can be retrained or an additional set of models can be trained and stored for the 25 new user. When the training data for training the speaker specific systems is not readily available, speaker independent systems tend to be used as well. In addition, as the number of users becomes large, storing a separate speaker specific speech model set for each user becomes 30 prohibitive in terms of memory requirements. Therefore, as

the number of users becomes large, speech recognition systems tend to be speaker independent.

- A common approach to improve the performance of speaker independent speech recognition systems is adaptation: adjusting either speech models or features in a manner appropriate to the current speaker and environment. A typical adaptation technique is model adaptation. Generally, speaker adaptation starts with speaker-independent speech models derived from one or more speakers and then, based on a small amount of speech from a new speaker, creates new speaker-dependent models so that the recognition of the new speaker is improved. For a more detailed explanation on model adaptation, the reader is invited to consult R. Schwartz and F Kubala, *Hidden Markov Models and Speaker Adaptation*, Speech Recognition and Understanding: Recent Advances, Eds: P. Laface et R. De Mori, Springer-Verlag, 1992; L. Neumeyer, A. Sankar and V. Digalakis, *A Comparative Study of Speaker Adaptation Techniques*, Proc. Of EuroSpeech'95, pp.1127-1130, 1995; J.-L. Gauvain, G.-H. Lee, *Maximum a Posteriori Estimation for Multivariate Gaussian Mixture Observations of Markov Chains*, IEEE. Trans. on Speech and Audio Processing, Vol.2, April 1994, pp.291-298; and C. J. Leggetter, P.C. Woodland, *Maximum Likelihood Linear Regression for Speaker Adaptation of Continuous Density Hidden Markov Models*, Computer, Speech and Language, Vol.9, 1995, pp171-185. The content of these documents is hereby incorporated by reference.
- A deficiency in the above-described methods is that they require a relatively large amount of data on the basis

of the usual relationship between training data and parametric complexity. While humans seem to be able to adapt to a new speech environment in just a few syllables, such speech recognition system adaptation requires considerably more adaptation data which may not be available.

- A common approach makes use of a single set of transformation parameters for a set of speech models. In this manner a reduced number of transformation parameters permits effective adaptation with a reduced amount of data. A deficiency in tying the transformation parameters is that any model specific transformation will not be reflected in the single set of parameters.
- Another method is an adaptation method described in Kuhn R. et al. (1998), "Eigenvoices for speaker adaptation," Proc. ICSLP '98, vol. 5, Sydney, pp.1771-1774. The content of this document is hereby incorporated by reference. This adaptation method requires less adaptation data than the methods mentioned previously. A deficiency of the method presented by Kuhn R. et al. (1998) is that it provides an improved performance for speakers for whom training data is available. A deficiency of this approach is that the acoustic characteristics of specific speakers, not part of the training data, is generally not captured by this method.

Consequently, there is a need in the industry for providing a method and apparatus for providing a speech recognition system capable of adapting speech models on the basis of a minimal amount of data, preferably from the first 5 word given by a given speaker.

Summary of the Invention

In accordance with a broad aspect, the invention 10 provides a speech recognition system comprising an input for receiving an input signal derived from a spoken utterance that is indicative of at least one speech element. The speech recognition system further comprises a first processing unit coupled to the input operative for 15 processing the input signal to derive from a speech recognition dictionary at least one speech model associated to a given speech element that constitutes a potential match to the at least one speech element. The speech recognition system further comprises a second processing unit coupled to 20 the first processing unit for generating a modified version of the at least one speech model on the basis of the input signal. The speech recognition system further comprises a third processing unit coupled to the second processing unit for processing the input signal on the basis of the modified 25 version of the at least one speech model. The third processing unit is operative to generate a recognition result indicative of whether the modified version of the at least one speech model constitutes a match to the input signal. The speech recognition system further provides an 30 output for releasing a signal indicative of the recognition result.

Advantageously, by allowing the at least one speech model to be modified on the basis of the recognition attempt and then effecting speech recognition on the basis of the
5 modified at least one speech model, adaptation of the speech modes is effected during the recognition process.

In a specific example of implementation, the first processing unit implements a first-pass recognition process
10 making use of low-complexity low-definition models to derive the speech model that constitutes a match to the input signal. The second processing unit implements an apparatus for generating a hybrid speech model by processing the speech model that constitutes a potential match to the input
15 signal to generate a modified version of the speech model. As a variant the second processing unit implements an apparatus for generating a complex speech model by processing the speech model that constitutes a potential match to the input signal to generate a modified version of
20 the speech model. The third processing unit comprises a re-scoring unit operative to process the modified version of the speech model and the input signal.

Preferably, the model groups are shared across
25 speakers. Advantageously, by allowing the model groups to be shared across speakers, separate speaker dependent model sets for each individual speaker need not be stored thereby resulting in a reduction in memory storage requirements.

30 In accordance with another broad aspect, the invention provides a method and apparatus for generating a hybrid

speech model associated to a given speech element. The hybrid speech model is suitable for use by a speech recognition device during speech recognition. The apparatus comprises an input for receiving an input signal derived
5 from a spoken utterance that contains at least one speech element that potentially matches the given speech element. The apparatus further comprises a model group associated to the given speech element comprising a plurality of speech models. Each speech model in the plurality of speech models
10 is a different representation of the given speech element. The apparatus further comprises a processing unit coupled to the input for processing the input signal and the model group to generate a hybrid speech model associated to the given speech element. The hybrid speech model is a
15 combination of speech models of the plurality of speech models, the combination being effected on the basis of the input signal derived from the spoken utterance. The apparatus further comprises an output for releasing a signal conveying the hybrid speech model associated to the given
20 speech element in a format suitable for use by a speech recognition device.

In a specific example of implementation, the hybrid speech model is derived by computing a linear combination of
25 the speech models in the model group.

In accordance with another broad aspect, the invention provides a computer readable medium comprising a program element suitable for execution by a computing apparatus for
30 implementing the above-described apparatus.

In accordance with another aspect, the invention provides an apparatus and method for generating a complex speech model associated to a given speech element suitable for use by a speech recognition device during speech

5 recognition. The apparatus comprises an input for receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element. The apparatus further comprises a model group associated to the given speech

10 element. The model group comprises a plurality of speech models, each speech model being a different representation of the given speech element. The model group comprises two sets of speech models namely a first set having speech models of a first type and a second set having speech models

15 of a second type. Each speech model of a first type in the first set is associated to a speech model of the second type in the second set. The apparatus further comprises a processing unit coupled to the input for processing the input signal and the model group to generate a hybrid speech

20 model associated to the given speech element. The hybrid speech model is a combination of speech models of the first type of the plurality of speech models, the combination being effected on a basis of the input signal derived from the spoken utterance. The apparatus further comprises

25 processing the hybrid speech model to generate a complex speech model associated to the given speech element, the complex speech model being a combination of speech models of the second type of the plurality of speech models. The apparatus further comprises an output for releasing a signal

30 indicative of the complex speech model associated to the

given speech element in a format suitable for use by a speech recognition device.

In a specific example of implementation, any speech model of a second type is indicative of a speech model having a higher complexity than a speech model of a first type to which it is associated.

Continuing with the specific example of implementation, the hybrid speech model is derived by computing a first linear combination of the speech models of the first type. The first linear combination is characterized by a first set of parameters indicative of weights associated to speech models of the first type. The complex speech model is derived by computing a second linear combination of the speech models of the second type. The second linear combination is characterized by a second set of parameters indicative of weights associated to speech models of the second type. The first set of parameters and the second set of parameters are indicative of substantially same weights.

In accordance with another broad aspect, the invention further provides a computer readable medium comprising a data structure for storing a plurality of model groups, each model group being associated to a respective speech element in a phonetic alphabet. Each model group comprises a plurality of speech models, each model group being suitable for use by a speech-processing device.

In a specific example of implementation, each model group comprises two sets of speech models namely a first set

having a set of speech models of a first type and a second set having speech models of a second type. Each speech model of a first type in the first set is associated to a speech model of the second type in the second set. Each speech

- 5 model of the second type is indicative of a speech model having a higher complexity than a speech model of the first type to which the speech model of the second type is associated.

10 For the purpose of this specification, the expression "speech element" is used to designate a basic speech recognition unit. In a specific example, a speech element is indicative of a data element selected from the set consisting of phones, diphones, syllables and words. It will

- 15 be readily apparent to the person skilled in the art that representations of phones such as allophones and triphones may be used without detracting from the spirit of the invention.

20 For the purpose of this specification, the expressions "model" and "speech model" are used to designate a mathematical representation of the acoustic properties of a speech element. Modeling speech elements is well-known in the art to which this invention pertains. Commonly used

- 25 models include Hidden Markov Models (HMMs) where each speech element is represented by a sequence of states and transitions between the states. Speech models may have different levels of complexity depending on the number of transitions, states and parameters characterizing the states
30 that are used to model the range of acoustic properties. The expressions "model" and "speech models" should be given a

broad interpretation to include an electronic representation of the speech models themselves, a symbol being associated to a respective speech model, a pointer to a memory location allowing to extract a speech model or any other
5 representation allowing a speech model to be extracted.

Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific
10 embodiments of the invention in conjunction with the accompanying figures.

Brief Description of the Drawings

15 Figure 1 shows a block diagram of an apparatus for generating a hybrid speech model in accordance with an embodiment of the invention;

Figure 2 shows a flow diagram of a method for generating a hybrid speech model in accordance with an embodiment of
20 the invention;

Figure 3 shows a computing apparatus including a program element for generating a hybrid speech model in accordance with an embodiment of the invention;

Figure 4 shows an apparatus for generating a complex speech
25 model in accordance with an embodiment of the invention;

Figure 5 shows a speech recognition system in accordance with an embodiment of the invention;

Figure 6 shows a high-level block diagram of an apparatus
30 for generating a speech model group in accordance with another aspect of the invention;

Figure 7 shows a flow diagram of a method for generating a speech model group in accordance with the spirit of the invention;

5 Figure 8 shows an apparatus for generating a speech model group in accordance with the spirit of the invention.

Detailed Description

In a first form of implementation, as shown in figure
10 1, the invention provides an apparatus for generating a hybrid speech model associated to a given speech element, the hybrid speech model being suitable for use by a speech recognition device during speech recognition. The apparatus comprises an input 100, a model group 102, a processing unit
15 104 and an output 106.

The input 100 is for receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element.
20 In a specific example of implementation, the input signal derived from a spoken utterance is indicative of a speaker specific speech model associated to the at least one speech element. The hybrid speech model is derived based on the hypothesis that the input signal is a match to the given
25 speech model.

The output 106 is for releasing a signal indicative of the hybrid speech model associated to the given speech element in a format suitable for use by a speech recognition
30 device.

The model group 102 associated to the given speech element comprises a plurality of speech models, each speech model being a different representation of the given speech element. In a specific example of implementation, each speech models is indicative of allophonic models having three states, each state having an eight component mixture of Gaussians. Other formats of speech models may be used without detracting from the spirit of the invention and are well known in the art to which this invention pertains.

10

As a variant, the apparatus comprises a database of model groups 110 comprising a plurality of model groups. Each model group is associated to a respective speech element and comprises a set of speech models. In a specific example of implementation, the database of model groups is stored in a data structure, such as a table, and comprises a plurality of records, each record being associated to a given speech element and comprising a set of speech models. In a very specific example of implementation the speech element is an element of a phonetic alphabet. The table below shows a set of records of the speech elements /a/, /s/ and /r/. The number of speech models in each model group may vary for each speech element without detracting from the spirit of the invention.

25

Speech Element	Speech model #1	Speech model #2	Speech model #3
/a/	MODEL_A/_a/	MODEL_B/_a/	MODEL_C/_a/
/s/	MODEL_A/_s/		
/r/	MODEL_A/_r/	MODEL_B/_r/	

In a specific implementation of the same variant, the input 100 is a first input and the input signal received at the first input 100 is a first input signal. The apparatus further comprises a second input 108 for receiving a second

- 5 input signal indicative of a data element identifying the given speech element. The processing unit 104 is operative to extract from the database of model groups 110 a certain model group associated to the data element received at the second input 108 identifying the given speech element.

10

The processing unit 104 is coupled to the input 100 for processing the input signal and the model group 102 to generate a hybrid speech model associated to the given speech element. The hybrid speech model is a combination of

- 15 speech models of the model group 102 on the basis of the input signal derived from the spoken utterance received at the input 100. The speech models in the model group define a space indicative of the possible values that a hybrid speech model associated to a given speech element may take on.
- 20 Preferably, the combination of speech models of the speech model group is such that the hybrid speech model is weighted toward the speaker specific speech model defined by the input signal received at the input 100.

25

In a specific example of implementation, the hybrid speech model is derived by computing a linear combination of the speech models in the model group 102. The linear combination is characterized by a set of parameters indicative of weights associated to speech models in the

30 model group 102. The combination is effected by selecting a set of weights that minimizes the difference between the

linear combination of speech models from the model group and the signal received at input 100 as well as constraining the hybrid model to lie within the space defined by the speech models in the model group. Preferably, the space defined by
5 the speech models in the model group 102 is a convex space.

The set weights are the solution to a convex optimisation problem. Mathematically, the weights are computed by minimising the following quadratic function:
10

Equation 1

$$J([\text{weight vector}]^T) = \| [m_1 \ m_2 \ m_3 \ \dots \ m_L] [\text{weight vector}]^T - m \| ^2$$

where m_1, m_2, \dots, m_L are speech models in the model group 102,

15 m is a speaker specific speech model derived from the signal received at input 100 and $[\text{weight vector}]^T$ is a vector of L entries comprising of a set of weights, each weight in the set of weights being assigned to a respective speech model in the model group. Equation 1 is minimised under the
20 convex-set constraints defined below:

- a) $\text{SUM}(\text{weight}_l; \text{ for } l=1 \text{ to } l=L) = 1$
- b) $\text{weight}_l > 0 \text{ for } 1 \leq l \leq L$

25 In this manner, the hybrid model defined by the combination:

Equation 2

$$\text{Hybrid model} = [m_1 \ m_2 \ m_3 \ \dots \ m_L] [\text{weight vector}]^T$$

30 will lie in the space defined by the speech models in the model group 102. Minimising a function of the type

described in equation 1 under defined constraints is well known in the art of numerical analysis and will not be described further here.

- 5 As shown in figure 2, the invention provides a method for generating a hybrid speech model associated to a given speech element suitable for use by a speech recognition device during speech recognition. The method comprises receiving 200 an input signal derived from a spoken
10 utterance that contains at least one speech element that potentially matches the given speech element. The method further comprises providing 202 a model group associated to the given speech element, said model group comprising a plurality of speech models. Each speech model in the
15 plurality of speech models is a different representation of the given speech element. The method further comprises processing 204 the input signal received at step 200 and the model group to generate a hybrid speech model associated to the given speech element. The hybrid speech model is a
20 combination of speech models in the plurality of speech models. The combination is effected on the basis of the input signal derived from the spoken utterance. The method further comprises releasing 206 a signal indicative of the hybrid speech model associated to the given speech element
25 in a format suitable for use by a speech recognition device.

- In a specific example of implementation, the apparatus depicted in figure 1 is implemented on a digital computing apparatus of the type depicted in figure 3 of the drawings.
30 Such a computing apparatus comprises a processor 300 coupled to a memory unit 302, the memory unit comprising a program

memory 304 portion storing a program element for execution by the processor for implementing the processing unit 104. The memory unit 302 further comprises a data memory 306 portion for storing the model group described in connection

- 5 with block 102 and optionally the database of model groups 110. As a variant, the processing unit 104 is a hardware device such as a ROM or other chip programmed on the basis of the above-described method.

- 10 In accordance with a second aspect, as shown in figure 4 of the drawings, the invention provides an apparatus for generating a complex speech model associated to a given speech element, the complex speech model being suitable for use by a speech recognition device during speech
15 recognition. The apparatus comprises an input 400, a model group 402, a processing unit 404 and an output 406.

- The input 400 is for receiving an input signal derived from a spoken utterance that contains at least one speech
20 element that potentially matches the given speech element. In a specific example of implementation, the input signal derived from a spoken utterance is indicative of a speaker specific speech model associated to the at least one speech element.

25

The output 406 is for releasing a signal indicative of the complex speech model associated to the given speech element in a format suitable for use by a speech recognition device.

30

The model group 402 associated to the given speech element comprises a plurality of speech models, each speech model being a different representation of the given speech element. The model group comprises two sets of speech models

- 5 namely a first set having speech models of a first type and a second set having speech models of a second type. Each speech model of a first type in the first set is associated to a speech model of the second type in the second set. In a specific example of implementation, any speech model of a
10 second type is indicative of a speech model having a higher complexity than a speech model of a first type to which it is associated.

In a specific example of implementation, each speech
15 model of a first type is indicative of a speech model having three states, each state having a single Gaussian as its model of observation density. Each speech model of a second type is indicative of a speech model having three states, each state having a eight component mixture of Gaussians.
20 Other formats of speech models may be used without detracting from the spirit of the invention and are well known in the art to which this invention pertains.

As a variant, the apparatus comprises a database of
25 model groups 410 comprising a plurality of model groups. Each model group is associated to a respective speech element and comprises a plurality of speech models, each model group being in format similar to that described in connection with block 402. In a specific example of
30 implementation, the database of model groups 410 is stored on a computer readable medium in a data structure, such as a

table, comprising a plurality of records. Each record is associated to a given speech element and comprises two sets of speech models namely a first set having speech models of a first type and a second set having speech models of a

- 5 second type. Each speech model of a first type in the first set is associated to a speech model of the second type in the second set. In a very specific example of implementation the speech element is an element of a phonetic alphabet. The table below shows a set of records of the speech elements /a/, /s/ and /r/. The number of speech models in each model group may vary for each speech element without detracting from the spirit of the invention.

Speech Element	Speech model #1	Speech model #2	Speech model #3
/a/	MODEL_A/_a/ COMPLEX_A/_a/	MODEL_B/_a/ COMPLEX_B/_a/	MODEL_C/_a/ COMPLEX_C/_a/
/s/	MODEL_A/_s/ COMPLEX_A/_s/		
/r/	MODEL_A/_r/ COMPLEX_A/_r/	MODEL_B/_r/ COMPLEX_B/_r/	

- 15 For example, in the above table the first set for speech element /a/ comprises {MODEL_A/_a/; MODEL_B/_a/; MODEL_C/_a/} and the second set for speech element /a/ comprises {COMPLEX_A/_a/; COMPLEX_B/_a/; COMPLEX_B/_a/}.

- 20 In a specific implementation of the same variant, the input 400 is a first input and the input signal received at the first input 400 is a first input signal. The apparatus further comprises a second input 408 for receiving a second

input signal indicative of a data element identifying the given speech element. The processing unit 404 is operative to extract from the database of model groups 410 a certain model group associated to the data element received at the 5 second input 408 identifying the given speech element.

The processing unit 404 is coupled to the input 400 for processing the input signal and the model group 402 to generate a hybrid speech model associated to the given 10 speech element. The hybrid speech model is a combination of speech models of the first type of the model group 402 on the basis of the input signal derived from the spoken utterance received at the input 400. The hybrid speech model is generated substantially in the same manner as 15 described previously in connection with figure 1 of the drawings.

In a specific example of implementation, the hybrid speech model is derived by computing a linear combination of 20 the speech models in the model group 402. The linear combination is characterized by a first set of parameters indicative of weights associated to speech models in the model group 402. As described previously, the hybrid model is expressed by as equation 2, reproduced below for the 25 reader's convenience:

Equation 2

$$\text{Hybrid model} = [m_1 \ m_2 \ m_3 \ \dots \ m_L] [\text{weight vector}]^T$$

30 where m_1, m_2, \dots, m_L are speech models in the model group 102, and $[\text{weight vector}]^T$ is a vector of L entries comprising of

a set of weights, each weight in the set of weights being assigned to a respective speech model in the model group and where the [weight vector]^T satisfies the following conditions:

5

- a) SUM(weight_l; for l=1 to l=L) =1
- b) weight_l > 0 for 1≤l≤L

The processing unit 404 is further operative to process
10 the hybrid speech model to generate a complex speech model associated to the given speech element. The complex speech model is a combination of speech models of the second type in the plurality of speech models.

15 In a specific example of implementation, the hybrid model is a first linear combination and is characterized by a first set of parameters indicative of weights associated to speech models of the first type. The complex speech model is derived by computing a second linear combination of
20 the speech models of the second type, the second linear combination being characterized by a second set of parameters indicative of weights associated to speech models of the second type. The first set of parameters and the second set of parameters are indicative of substantially
25 same weights.

Mathematically, the complex speech model can be expressed as follows:

30

Equation 3

$$\text{Complex model} = [b_1 \ b_2 \ b_3 \dots \ b_L] [\text{weight vector}]^T$$

where b_1, b_2, \dots, b_L are speech models of the second type in the model group 402 and $[weight\ vector]^T$ is the vector of L entries comprising of a set of weights derived from the 5 calculations of the hybrid model.

As a variant, the processing unit 404 is further operative to implement a fall-back unit for an input signal derived from a spoken utterance that cannot be represented 10 accurately by the plurality of speech models in the model group 402.

In a specific example of implementation, the complex speech model depends on how close the hybrid model can be to 15 the input signal received at input 400. Mathematically, this was expressed as the minimum value that $J([weight\ vector]^T)$ can take on in equation 1 reproduced below for the user's convenience.

20

Equation 1

$$J([weight\ vector]^T) = ||[m_1\ m_2\ m_3\ \dots\ m_L][weight\ vector]^T - m||^2$$

If $J([weight\ vector]^T)$ is very small, then the complex model is a good approximation of the input speech signal.

25 If $J([weight\ vector]^T)$ is large then the approximation is does not provide a good approximation.

In a specific example, the fall-back unit is operative to condition the linear combination of the speech models of 30 the second type on the basis of the how close the hybrid model can be to the input signal received at input 400.

This is effected by providing an interpolation between a base complex speech models (b_0) and the linear combination of the speech models of the second type shown in equation 3. In this specific example, the base complex speech models

- 5 (b_0) is a speaker independent speech model associated to the speech element. The base model b_0 may be stored in the model group 402 or may be received as a parameter by the apparatus depicted in figure 4 without detracting from the spirit of the invention.

10

A mathematical example will better illustrate this concept. The example below is shown for the purpose of example only and the person skilled in the art will readily find that variations are possible without detracting from

- 15 the spirit of the invention. The computation of the complex model is modified to account for how well the hybrid model approximated the input signal. Equation 3 above is modified as follows:

20

Equation 3 (modified)

Complex model = $f^* [b_1 \ b_2 \ b_3 \ ... \ b_L] [\text{weight vector}]^T + (1-f)b_0$

For the present example, the following relation is used:

25

$$f = \min[J([\text{weight vector}]^T) / ||m||^2, 1]$$

- When the hybrid model is a good approximation of the input signal, the $J([\text{weight vector}]^T)$ is small relative to the 30 speaker specific model m and f will be close to unity. On the other hand, if the model group 402 cannot represent the

given least one speech element, the $\mathcal{T}([\text{weight vector}]^T)$ can be quite large, in which event the model adaptation falls back to the base model b_0 .

5 In accordance with another aspect, the invention provides a speech recognition system of the type depicted in figure 5 of the drawings. The speech recognition system comprises an input 510, a first processing unit 500, a second processing unit 506, a third processing unit 508 and
10 an output 512.

The input 510 is for receiving an input signal derived from a spoken utterance that is indicative of at least one speech element.

15 The first processing unit 500 is coupled to the input 510 and is operative for processing the input signal to derive from a speech recognition dictionary 502 a speech model associated to a given speech element that constitutes
20 a potential match to the at least one speech element. In a specific example of implementation, the first processing unit 500 is operative to perform rough probabilistic calculations and to extract from the speech recognition dictionary 502 a list of possible candidate speech models
25 for the spoken utterance received at the input 510. The first processing unit allows the system to eliminate the entries in the speech recognition dictionary 502 that are most unlikely to constitute a match to the spoken utterance. Typical algorithms that can be used at this stage include
30 the fast score estimation and the graph search algorithms. As a reference, the reader is invited to consult Gupta V.

N., Lennig M., Mermelstein P. "A fast search strategy in a large vocabulary word recogniser INRS-Telecommunications. J. Acoust. Soc. Am. 84 (6), December 1988, p.2007 and U.S. patent #5,515,475 by inventors Gupta V. N. & Lennig M. The
5 content of these documents is incorporated herein by reference. Other algorithms may be used here without detracting from the spirit of the invention. In a specific example of implementation, the first processing unit 500 is operative for generating a speaker specific speech model
10 derived on the basis of the input spoken utterance received at input 510, the speaker specific speech model being indicative of the acoustic characteristics of the least one speech element. Continuing with this specific example, the speaker specific speech model is derived by conditioning the
15 at least one speech model using a maximum *a posteriori* adaptation method. For more information on using maximum *a posteriori* (MAP) adaptation, the reader is invited to consult Gauvain et al. (1994), "maximum *a posteriori* estimation for multivariate Gaussian mixture observations of
20 Markov chains", IEEE Transaction Speech Audio Process. 2, pages 291 to 298.

The second processing unit 506 is coupled to the first processing unit 500. The second processing unit is operative
25 to generate a modified version of the speech model derived from the dictionary by conditioning the latter on the basis of the input spoken utterance received at input 510. In other words, the modified version of the speech model derived from the dictionary is derived based on the
30 hypothesis that the input spoken utterance is a match to the speech model derived from the dictionary. By allowing the

speech model derived from the dictionary to be modified for each spoken attempt, adaptation of the recognition process is permitted from the first spoken utterance. In addition, separate speaker dependent model sets for each individual

- 5 speaker need not be stored since model groups can be shared across speakers thereby resulting in a reduction in memory storage requirements. This is particularly useful for a speech recognition system operates in an environment with a large number of users. For example, this is particularly
10 useful in the context of an automatic directory assistance system making use of the speech recognition system.

In a first form of implementation, the second processing unit 506 is implemented substantially as

- 15 processing unit 104 described previously. The processing unit is operative to generate a hybrid speech model associated to the speech model derived from the dictionary on the basis of a signal derived from the input spoken utterance. The second processing unit 506 is coupled to the
20 first processing unit 500 via a first coupling member 514 and a second coupling member 520. In this first form of implementation, the first coupling member corresponds to the first input 100 and is for receiving the speaker specific speech model generated by the first processing unit 500.
25 The second coupling member 520 corresponds to the second input 108 and is for receiving the speech model derived by the first processing unit 500 from the speech recognition dictionary 502. The second processing unit 506 is further coupled to a database of model groups 510 via a third
30 coupling link 518. The database of model groups 510 is implemented in a matter similar to that described in

connection with block 110 in figure 1 and comprises a plurality of models groups 502 implementation in a matter similar to that described in connection with block 102. The second processing unit 506 releases a signal indicative of a
5 hybrid speech model at an output coupling member 516 corresponding to output 106 in figure 1.

In a second form of implementation, the second processing unit 506 is implemented substantially as
10 processing unit 404 described previously. The processing unit is operative to generate a complex speech model associated to the speech model derived from the dictionary on the basis of a signal derived from the input spoken utterance. The second processing unit 506 is coupled to the
15 first processing unit 500 via a first coupling member 514 and a second coupling member 520. In this second form of implementation, the first coupling member corresponds to the first input 400 and is for receiving the speaker specific speech model generated by the first processing unit 500.
20 The second coupling member 520 corresponds to the second input 408 and is for receiving the speech model derived by the first processing unit 500 from the speech recognition dictionary 502. The second processing unit 506 is further coupled to a database of model groups 510 via a third
25 coupling link 518. The database of model groups 510 is implemented in a matter similar to that described in connection with block 410 in figure 4 and comprises a plurality of models groups 502 implementation in a matter similar to that described in connection with block 402. The
30 second processing unit 506 releases a signal indicative of a

complex speech model at an output coupling member 516 corresponding to output 406 in figure 4.

The third processing unit 508 is coupled to the second
5 processing unit 506 via coupling link 516. The third processing unit is operative for processing the input spoken utterance received at input 510 on the basis of the modified version of the at least one speech model generated by the second processing unit 506. The third processing unit 508
10 generates a recognition result indicative of whether the modified version of the at least one speech model constitutes a match to the input signal. The output 512 is for releasing a signal indicative of the recognition result.

15 In a specific example of implementation, the third processing unit 516 is a re-scoring unit operative to perform more precise speech recognition calculations than the first processing unit 500. The speech recognition calculations are effected by performing comparisons between
20 the acoustic characteristics of the input spoken utterance received at input 510 and the modified version of the speech model generated by the second processing unit 506. The third processing unit 508 may make use of techniques such as the Viterbi algorithm with complete allophone models and model
25 distances. Although these require heavy complex computations, the number of candidates for which the computation must be performed has been greatly reduced. The top scoring recognition candidate is then released as the recognition result at output 512.

In accordance with another aspect, as shown in figure 6, the invention provides an apparatus 606 for generating a speech model group associated to a given speech element, the speech model group containing a plurality of speech models.

- 5 The apparatus 606 comprises an input 608, a processing unit 606 and an output 604.

The input 608 is for receiving a plurality of data items derived from a plurality of spoken utterances, each

- 10 utterance in the plurality of spoken utterances being associated to the given speech element. In a specific example of implementation, this is effected by providing a computer readable storage medium coupled to the apparatus 606 via the input 608. The computer readable medium.
- 15 contains a plurality of data items associated to the given speech element. The data items may be in the form of speech tokens, in the form of feature vectors derived from speech tokens or in any other representation providing acoustical information on the speech tokens without detracting from the
- 20 spirit of the invention.

The processing unit 602 is coupled to the input 602 and is operative for generating a speech model group associated

- 25 to the given speech element comprising a plurality of speech models, each speech model being a different representation of the given speech element. In a specific example of implementation, the functionality of the processing unit 602 will be best understood in connection with the flow chart
- 30 depicted in figure 7.

The processing unit 602 is operative for receiving 700 a plurality of data items associated to a given speech element. The processing unit 602 processes each data item to extract acoustic characteristics. This is effected by 5 deriving 702 for each data item a data item specific speech model. In a specific example of implementation, the data item specific speech model is derived by applying a maximum *a posteriori* adaptation method on the basis of the speech token to a reference model associated to the given speech 10 element. In this specific example the reference speech model in a speaker independent model associated to the given speech element. The data items are then grouped 704 into a set of data item clusters, each data item cluster having at least one data item, on a basis of the similarity between the 15 acoustic characteristics of the data items. Clustering may be effected by any suitable clustering technique such as the K-means algorithm. A description of a K-means algorithm, also known as the LBG algorithm, is given in Makhoul et al. "Vector quantization in speech coding," Proceedings of the 20 IEEE, vol. 73, pp. 1551-1588, Nov 1985 whose content is hereby incorporated by reference. The reader is also invited to consult Rabiner et al. (1979) "Applications of clustering techniques to speaker-trained isolated word recognition," Bell System Technical Journal, vol. 58, pp.2217-2231 whose 25 content is hereby incorporated by reference. Such clustering techniques are well known in the art to which this invention pertains and will not be further described here. Following this the processing unit 602 is further operative to process 706 each data item cluster to generate a cluster specific 30 model. In a specific example of implementation, a given cluster specific speech model is derived by applying a

maximum *a posteriori* adaptation method on the basis of the data items in the associated data item cluster to a reference model associated to the given speech element. The signal indicative of the set of cluster specific models is

- 5 then released 708 at output 604. The set of cluster specific models forms a speech model group comprising a plurality of speech models, each speech model in the speech model group being associated to a respective data item cluster. Each speech model in the speech model group is an
10 approximation of acoustic properties of data items in the respective data item cluster.

Figure 8 shows a block diagram of an apparatus for generating the speech model group associated to a given

- 15 speech element in accordance with the spirit of the invention. The apparatus comprises a speaker independent model unit comprising a data element indicative of a speaker independent model associated to the given speech element. The apparatus further comprises a set of data item clusters
20 802 804 806 808, each data item cluster containing a plurality of data items. The apparatus further comprises at least one adaptation units 810 for processing the speaker independent speech model on the basis of a given data item cluster to derive a cluster specific model. The adaptation
25 unit may be a single module operative to process each data item clusters 802 804 806 808 to derive the respective cluster specific models 812 814 816 818 or a set of adaptation units without detracting from the spirit of the invention. The adaptation unit 810 may make use of any
30 suitable adaptation method such as maximum *a posteriori*

(MAP) adaptation and Maximum Likelihood Linear Regression (MLLR).

The above-described apparatus for generating a speech model group described in the specification can also be implemented on any suitable computing platform. Such a computing platform typically includes a processor and a memory or computer readable medium connected to the processor by a data communication bus. The memory stores the data and the instructions of the program element implementing the functional blocks depicted in the drawings and described in the specification. In a specific example, the program element implements the processing unit 602. In another specific example, the program element is operative for implementing the adaptation unit 810. In a specific example, the plurality of data items associated to a given speech element is stored in the data portion of the memory. The program element operates on the data 904 in accordance with the algorithms described above to generate a speech model group using the techniques described in this specification.

Although the present invention has been described in considerable detail with reference to certain preferred embodiments thereof, variations and refinements are possible without departing from the spirit of the invention. Therefore, the scope of the invention should be limited only by the appended claims and their equivalents.

Claims:

1. An apparatus for generating a hybrid speech model associated to a given speech element, said hybrid speech model being suitable for use by a speech recognition device during speech recognition, said apparatus comprising:

- an input for receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element;
- a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element;
- a processing unit coupled to the input for processing the input signal and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models in said plurality of speech models on the basis of the input signal derived from the spoken utterance;
- an output for releasing a signal indicative of said hybrid speech model associated to the given speech element in a format suitable for use by a speech recognition device.

- RECEIVED
U.S. PATENT AND TRADEMARK OFFICE
JULY 1 1997
10693RMUS01U
2. An apparatus as defined in claim 1, wherein the given speech element is an element selected from the group consisting of phones, diphones, syllables and words.
- 5 3. An apparatus as defined in claim 2, wherein said input signal derived from a spoken utterance is indicative of a speaker specific speech model associated to the at least one speech element.
- 10 4. An apparatus as defined in claim 3, wherein said hybrid speech model is weighted toward the speaker specific speech model.
- 15 5. An apparatus as defined in claim 4, wherein said hybrid speech model is derived by computing a linear combination of the speech models in said model group.
- 20 6. An apparatus as defined in claim 5, wherein said input is a first input and wherein said input signal is a first input signal, said apparatus further comprising:
- 25 a) a second input for receiving a second input signal conveying a data element identifying the given speech element;
- 30 b) a database of model groups comprising a plurality of model groups, each model group being associated to a respective speech element, each model group comprising a set of speech models; said processing unit being further operative for extracting from said database of model groups a certain model group associated to the data element

received at said second input identifying the given speech element.

7. A method for generating a hybrid speech model
5 associated to a given speech element, said hybrid speech model being suitable for use by a speech recognition device during speech recognition, said method comprising:

- receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element;
- providing a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element;
- processing the input signal and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models in said plurality of speech models on the basis of the input signal derived from the spoken utterance;
- releasing a signal indicative of said hybrid speech model associated to the given speech element in a format suitable for use by a speech recognition device.

8. A computer readable medium comprising a program element suitable for execution by a computing apparatus for generating a hybrid speech model associated to a given speech element, said hybrid speech model being suitable for use by a speech recognition device during speech recognition, said computing apparatus comprising:

- 5 - a memory unit for storing a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element;
- 10 - a processor operatively connected to said memory unit, said program element when executing on said processor being operative for implementing:
- 15 - an input for receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element;
- 20 - a processing unit coupled to the input for processing the input signal and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models in said plurality of speech models on the basis of the input signal derived from the spoken utterance;
- 25 - an output for releasing a signal indicative of said hybrid speech model associated to the given speech element in a format suitable for use by a speech recognition device.
- 30

9. An apparatus for generating a complex speech model associated to a given speech element, said complex speech model being suitable for use by a speech recognition device during speech recognition, said apparatus comprising:

- an input for receiving an input signal derived from a spoken utterance that contains at least one speech element that potentially matches the given speech element;
- a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element, said model group comprising two sets of speech models namely a first set having speech models of a first type and a second set having speech models of a second type, each speech model of a first type in said first set being associated to a speech model of the second type in the second set;
- a processing unit coupled to the input for:
 - a) processing the input signal and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models of the first type in said plurality of speech models on the basis of the input signal derived from the spoken utterance;

5 b) processing the hybrid speech model to generate a complex speech model associated to the given speech element, said complex speech model being a combination of speech models of the second type in said plurality of speech models;

10 - an output for releasing a signal indicative of said complex speech model associated to the given speech element in a format suitable for use by a speech recognition device.

15 10. An apparatus as defined in claim 9, wherein any speech model of a second type is indicative of a speech model having a higher complexity than a speech model of a first type to which it is associated.

20 11. An apparatus as defined in claim 10, wherein the given speech element is indicative of a data element selected from the set consisting of phones, diphones, syllables and words.

25 12. An apparatus as defined in claim 11, wherein said input signal derived from a spoken utterance is indicative of a speaker specific speech model associated to the at least one speech element.

30 13. An apparatus as defined in claim 12, wherein said hybrid speech model is weighted toward the speaker specific speech model.

14. An apparatus as defined in claim 13, wherein said hybrid speech model is derived by computing a linear combination of the speech models of the first type.

5 15. An apparatus as defined in claim 14, wherein said linear combination is a first linear combination and is characterized by a first set of parameters indicative of weights associated to speech models of the first type, said complex speech model being derived by computing a
10 second linear combination of the speech models of the second type, said second linear combination being characterized by a second set of parameters indicative of weights associated to speech models of the second type.

15 16. An apparatus as defined in claim 15, wherein said first set of parameters and said second set of parameters are indicative of substantially same weights.

17. An apparatus as defined in claim 10, wherein said input
20 is a first input and wherein said input signal is a first input signal, said apparatus further comprising:

a) a second input for receiving a second input signal indicative of a data element identifying the given speech element;

25 b) a database of model groups comprising a plurality of model groups, each model group being associated to a respective speech element, each model group comprising two sets of speech models namely a first set having speech models of a first type and a second set having speech models of a second type, each speech model of a first
30 type being associated to a respective speech element.

type in said first set being associated to a speech model of the second type in the second set;

5 said processing unit being further operative for extracting from said database of model groups a certain model group associated to the data element received at said second input identifying the given speech element.

10 18. A computer readable medium comprising a data structure for storing a plurality of model groups, each model group being associated to a respective speech element in a phonetic alphabet, each model group comprising a plurality of speech models, each model group being suitable for use by a processing device.

15 19. A computer readable medium as defined in claim 18, wherein the given speech element is indicative of a data element selected from the set consisting of phones, diphones, syllables and words.

20 20. A computer readable medium as defined in claim 19, wherein each model group comprises two sets of speech models namely a first set having a plurality of speech 25 models of a first type and a second set having a plurality of speech models of a second type, each speech model of a first type in said first set being associated to a speech model of the second type in the second set, each speech model of the second type being indicative of a speech model having a higher complexity 30

that a speech model of the first type to which the speech model of the second type is associated.

21. A speech recognition system comprising:

- 5 - an input for receiving an input signal indicative of a spoken utterance that is indicative of at least one speech element;
- 10 - a first processing unit coupled to said input operative for processing the input signal to derive from a speech recognition dictionary at least one speech model associated to a given speech element that constitutes a potential match to the at least one speech element;
- 15 - a second processing unit coupled to said first processing unit for generating a modified version of the at least one speech model on the basis of the input signal;
- 20 - a third processing unit coupled to said second processing unit for processing the input signal on the basis of the modified version of the at least one speech model to generate a recognition result indicative of whether the modified version of the at least one speech model constitutes a match to the input signal;
- 25 - an output for releasing a signal indicative of the recognition result.

22. A speech recognition system as defined in claim 21, wherein said first processing unit is operative for

- 30 generating a speaker specific speech model derived on the basis of the input signal, the speaker specific

speech model being indicative of the acoustic characteristics of the least one speech element.

23. A speech recognition system as defined in claim 22,

5 wherein said modified version of the at least one speech model is indicative of a hybrid speech model associated to the given speech element.

24. A speech recognition system as defined in claim 23,

10 wherein said second processing unit comprises:

- coupling member for allowing data exchange between the first processing unit and the second processing unit, said coupling member being suitable for receiving the speaker specific speech model derived from the input signal;
- a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element;
- a functional unit coupled to the coupling member for processing the speaker specific speech model and the model group to generate the hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models in said plurality of speech models on the basis of the speaker specific speech model;

5 - an output coupling member for allowing data exchange between the second processing unit and the third processing unit, said output coupling member being suitable for releasing a signal indicative of the hybrid speech model associated to the given speech element.

10 25. A speech recognition system as defined in claim 24, wherein said hybrid speech model is weighted toward the speaker specific speech model.

15 26. A speech recognition system as defined in claim 24, wherein said hybrid speech model is derived by computing a linear combination of the speech models in said group of speech models.

20 27. A speech recognition system as defined in claim 24, wherein said coupling member is a first coupling member, said second processing unit further comprising:

25 a) a second coupling member for allowing data exchange between the first processing unit and the second processing unit, said second coupling member being suitable for receiving a data element identifying the given speech element;

30 b) a database of model groups comprising a plurality of model groups, each model group being associated to a respective speech element, each model group comprising a set of speech models; said functional unit being further operative for extracting from said database of model groups a certain model group associated to the data element

received at said second coupling member identifying the given speech element.

28. A speech recognition system as defined in claim 22,

5 wherein said modified version of the at least one speech model is indicative of a complex speech model associated to the given speech element.

29. A speech recognition system as defined in claim 28,

10 wherein said second processing unit comprises:

- coupling member for allowing data exchange between the first processing unit and the second processing unit, said coupling member being suitable for receiving the speaker specific speech model derived from the input signal;

- a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element, said model group comprising two sets of speech models namely a first set having speech models of a first type and a second set having speech models of a second type, each speech model of a first type in said first set being associated to a speech model of the second type in the second set;

- a functional unit coupled to the coupling member for:

a) processing the speaker specific speech model and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models of the first type in said plurality of speech models on the basis of the speaker specific speech model;

b) processing the hybrid speech model to generate a complex speech model associated to the given speech element, said complex speech model being a combination of speech models of the second type in said plurality of speech models;

- an output coupling member for allowing data exchange between the second processing unit and the third processing unit, said coupling member being suitable for releasing a signal indicative of the complex speech model associated to the given speech element.

30. A speech recognition system as defined in claim 29, wherein any speech model of the second type is indicative of a speech model having a higher complexity than a speech model of the first type to which it is associated.

31. A speech recognition system as defined in claim 30, wherein said hybrid speech model is weighted toward the speaker specific speech model.

32. A speech recognition system as defined in claim 31,
wherein said hybrid speech model is derived by computing
a linear combination of the speech models of the first
5 type.

33. A speech recognition system as defined in claim 32,
wherein said linear combination is a first linear
combination and is characterized by a first set of
10 parameters indicative of weights associated to speech
models of the first type, said complex speech model being
derived by computing a second linear combination of the
speech models of the second type, said second linear
combination being characterized by a second set of
15 parameters indicative of weights associated to speech
models of the second type.

34. A speech recognition system as defined in claim 33,
wherein said first set of parameters and said second set
20 of parameters is indicative of substantially the same
weights.

35. An apparatus for generating a hybrid speech model
associated to a given speech element, said hybrid
25 speech model being suitable for use by a speech
recognition device during speech recognition, said
apparatus comprising:

30 - means for receiving an input signal derived
from a spoken utterance that contains at least
one speech element that potentially matches the
given speech element;

PCT/INDIA2008/000001

- a model group associated to the given speech element, said model group comprising a plurality of speech models, each speech model of said plurality of speech models being a different representation of the given speech element;
- means for processing the input signal and the model group to generate a hybrid speech model associated to the given speech element, said hybrid speech model being a combination of speech models in said plurality of speech models on the basis of the input signal derived from the spoken utterance;
- means for releasing a signal indicative of said hybrid speech model associated to the given speech element in a format suitable for use by a speech recognition device.

36. A speech recognition system comprising:

- means receiving an input signal indicative of a spoken utterance that is indicative of at least one speech element;
- means for processing the input signal to derive from a speech recognition dictionary at least one speech model associated to a given speech element that constitutes a potential match to the at least one speech element;
- means for generating a modified version of the at least one speech model on the basis of the input signal;

5

- means for processing the input signal on the basis of the modified version of the at least one speech model to generate a recognition result indicative of whether the modified version of the at least one speech model constitutes a match to the input signal;
- means for releasing a signal indicative of the recognition result.

ABSTRACT OF THE DISCLOSURE

A speech recognition system having an input for receiving an input signal indicative of a spoken utterance that is
5 indicative of at least one speech element. The system further includes a first processing unit operative for processing the input signal to derive from a speech recognition dictionary a speech model associated to a given speech element that constitutes a potential match to the at
10 least one speech element. The system further comprised a second processing unit for generating a modified version of the speech model on the basis of the input signal. The system further provides a third processing unit for processing the input signal on the basis of the modified
15 version of the speech model to generate a recognition result indicative of whether the modified version of the at least one speech model constitutes a match to the input signal. The second processing unit allows the speech model to be modified on the basis of the recognition attempt thereby
20 allowing speech recognition to be effected on the basis of the modified speech model. This permits adaptation of the speech models during the recognition process. The invention further provides an apparatus, method and computer readable medium for implementing the second processing unit.




FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7




FIGURE 8

Docket No.
85773-161

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND APPARATUS PROVIDING HYPOTHESIS DRIVEN SPEECH MODELLING FOR USE IN SPEECH RECOGNITION

the specification of which

(check one)

is attached hereto.

was filed on _____ as United States Application No. or PCT International Application Number _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)	Priority	Not Claimed
(Number)	(Country)	(Day/Month/Year Filed)
(Number)	(Country)	(Day/Month/Year Filed)
(Number)	(Country)	(Day/Month/Year Filed)

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional application(s) listed below:

(Application Serial No.)	(Filing Date)
--------------------------	---------------

(Application Serial No.)	(Filing Date)
--------------------------	---------------

(Application Serial No.)	(Filing Date)
--------------------------	---------------

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)
--------------------------	---------------	--

(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)
--------------------------	---------------	--

(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)
--------------------------	---------------	--

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (*list name and registration number*)

Raymond Trudeau,
Stephan P. Georgiev

Reg. no. 26,412;
Reg. no. 37,563

Send Correspondence to: **SMART & BIGGAR**
1000 de la Gauchetiere Street West
Suite 3400, Montreal, Quebec H3B 4WS
CANADA

Direct Telephone Calls to: *(name and telephone number)*
Stephan P. Georgiev at (514) 954-1500

Full name of sole or first inventor Stephen Douglas PETERS	Date <i>14/12/99</i>
Sole or first inventor's signature <i>[Signature]</i>	
Residence 82 Pt Claire Avenue, Pt Claire, Quebec H9S 4M5 CANADA	
Citizenship Canadian	
Post Office Address Same as above	

Full name of second inventor, if any Daniel BOIES	Date <i>14/12/99</i>
Second inventor's signature <i>[Signature]</i>	
Residence 1 Place de Guise #12, Candiac, Quebec J5R 4P7 CANADA	
Citizenship Canadian	
Post Office Address Same as above	

Full name of third inventor, if any Benoit DUMOULIN	Date
Third inventor's signature <i>Benoit Dumoulin</i>	17/12/99
Residence 4-2110 St Hubert, Montreal, Quebec H2L 3Z4 CANADA	
Citizenship Canadian	
Post Office Address Same as above	

Full name of fourth inventor, if any	
Fourth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	

Full name of fifth inventor, if any	
Fifth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	

Full name of sixth inventor, if any	
Sixth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	